Experimental and Theoretical Evaluation of Current and Proposed Markets Including Effects of Ancillary Services

Robert J. Thomas
Timothy Mount
Ray Zimmerman
Simon Ede

March 2000

Sponsored by
Consortium for Electric Reliability Technology Solutions
Office of Energy Efficiency and Renewable Energy
U.S. Department of Energy
Washington, DC 20585

Power Systems Engineering Research Center
Cornell University
Ithaca, NY 14853

1.0 Introduction

The original objective of this work was to use experimental economic and other methods to simulate the performance of both current and proposed market (including ancillary services) designs and to explore, test, and demonstrate both theoretical and experimental economic approaches that simulate market performance in ways that accurately reflect the physical capabilities and limitations of the electric power system and the risks inherent in linking it to volatile markets. We have explored the rules for the PJM market, the New England market and have devised a market of our own that has the property that, at least in terms of simulation results, mitigates the non-market-power price spike behavior seen in real markets.

The strong relationship between a properly designed ancillary services market and price spikes emanating from other than market power behavior is an important discovery produced as a result of this work. The fact that price spike behavior in markets such as New England may be a result of a poorly designed ancillary services market is significant.

This report first discusses the unique experimental platform POWERWEB and the LEEDR lab, and their roles in this project. PJM and New England market data were analyzed to determine their efficiency as well to try and quantify the effects of similiarities and differences between market rules. Finally, experiments were performed to try and determine if market characteristics were indeed captured. We are especially interested in ancillary service market design. An important result is that we are able to design autonomous agents to operate the markets, where these agents capture essential market characteristics. This means we can use these agents to examine the operation of systems.

2.0 The PowerWeb Experimental Platform and Past Experiments Related to This Work

In this section we introduce the POWERWEB platform being used to examine alternative ancillary service markets. To illustrate the effectiveness of POWERWEB and to highlight two important market characteristics, with ramifications for ancillary service markets, two previous series of experiments dealing with market power and the unit commitment problem will be discussed.

Cornell University is the site of the Laboratory of Experimental Economics and Decision Research (LEEDR) as well as one of the sites of the Power Systems Engineering Research Center (PSERC). LEEDR recruits students at Cornell to participate in a wide range of economic experiments. In addition, PSERC recruits its industrial sponsors to participate in the experiments to ensure that a strong correlation exists between student participant response and highly experienced professionals. The experiments are usually competitive by nature and the participants receive performance-related pay. Work relating to electricity markets has employed a web-based software platform called POWERWEB. POWERWEB was designed and built at Cornell as a PSERC project. It is designed to be a simulation environment for experimentally testing various formats for competitive day-ahead electric energy markets. Importantly, POWERWEB restricts the market to solutions that abide by physical laws of real and reactive power flow while respecting network limitations such as line constraints, voltage limits and generation real and reactive power limits.

There are two main advantages to laboratory methods in economics. Firstly, the experiments and results are replicable and so it is possible to verify the findings independently. Most information in the electricity market is private and proprietary. It is difficult, consequently, to gather sufficient information to verify conclusions. Secondly, laboratory conditions can be set to control for extraneous circumstances which would be difficult to avoid and hard (or expensive) to mitigate in the real world. This allows the researcher to eliminate "noise" and focus upon the underlying validity of the theory or policy at hand.

A cause of concern is the selection of subjects for the experiments and their similarity to the real world decision makers. Experiments conducted at LEEDR, however, have shown in these electrical power experiments that a student pool of subjects performed in a similar manner to trained electricity industry professionals. There is an added qualitative advantage to using "inexperienced" subjects. If those subjects confirm the validity of the theory, it is likely that experienced decision makers will more easily make the same decisions. Indeed, the difference in behavior, in experiments to test for market power, between the students and the professionals was the speed with which the subjects figured out that they had the ability to raise prices above marginal costs.

3.0 The POWERWEB Platform: An Overview

Because of operational constraints on a power system, it is necessary to have a central agent acting as an independent system operator (ISO). In the previous implementations of POWERWEB, the ISO received offers to sell power from independently owned generation facilities. Based on a forecasted demand profile for the next day and the information gathered from the generator's offers, the ISO computed the optimal generator set points along with a corresponding price schedule which will allow the system to meet changing demand while satisfying all operational constraints.

As a web-based tool, POWERWEB may be used in several capacities. It can be utilized in a tightly controlled setting where a well-defined group of subjects are used for a very specific set of market experiments. It can also be used in a more open environment in which anyone on the web can log in and "play" as a generator competing against other generators, controlled by other humans or computer algorithms (agents), to generate power profitably. In either case, since POWERWEB is web-based it is accessible at all times to anyone with proper authorization, as long as the servers are up and running.

3.1 A Typical Session

To eliminate the need to coordinate accesses (via phone, e-mail, etc.) and to prevent one user's actions from interfering with another's, all accesses occur in the context of a given "session". The session specifies which power system is being simulated, who "owns" which system resources (generators, etc.), and what market mechanism is in use. Multiple sessions can be active at any given time and activity in each is completely independent of the others. Typically, a user in a session will "own" one or more generating plants.
After logging in as a generator in a simple auction session, for instance, the user is taken to the *Offer Submission* page such as shown in Figure **1**, which displays the cost and capacity information for their generator. Here they can enter offers to sell power to the ISO.

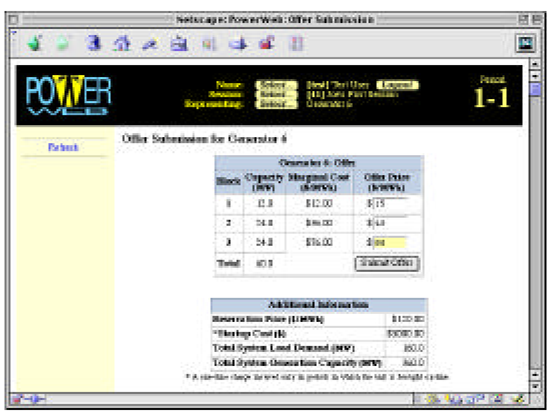

Figure 1: *Offer Submission* Page
When all participants have submitted their offers, POWERWEB's computational engine runs the auction according to the rules specified and reports back the results to the user. The *Auction Results* page is shown in Figure 2.

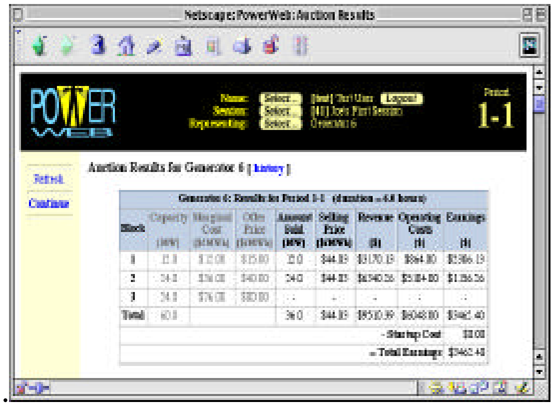

Figure 2: *Auction Results* Page

POWERWEB also has the capability to provide differing levels of information to the players, as specified by the experimenters. In a full information setting, each user would have access to the system information area, which gives tabular summaries of the system operation conditions as well as a "live" one-line diagram of the power system. Figure 3 shows the one-line diagram of a 6 generator, 30 bus system in POWERWEB's database. This diagram is generated dynamically by a Java applet from information retrieved from a relational database server. The diagram can be panned and zoomed and it is interactive in that clicking on an object such as a line, bus, generator, or load will query the database for information about the object. For example, selecting a bus will display the current information about real and reactive flows into and out of the bus as well as information about the current voltage level of the bus.

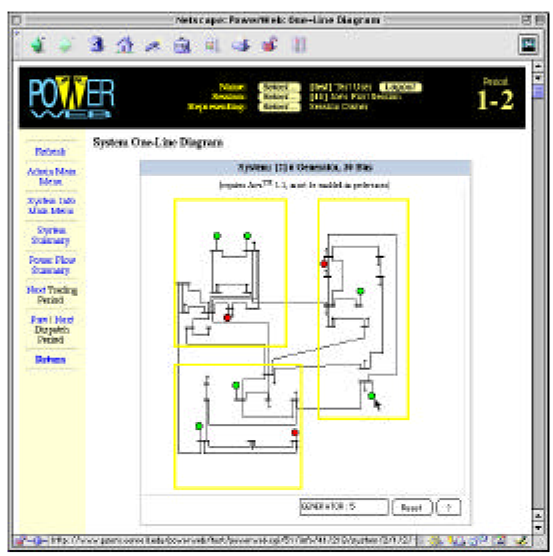

Figure 3: POWERWEB one-line diagram display, showing 30-bus system

The POWERWEB User's Manual, available from the POWERWEB home page at <http://www.pserc.cornell.edu/powerweb/> has more details regarding POWERWEB's functionality.

3.2 The Underlying Optimal Power Flow

At the heart of the POWERWEB computational engine is an optimal power flow (OPF) program that is executed by the ISO in response to offers submitted in an auction. The market activity rules determine what offers are valid, but it is the ISO's role to ensure the safe and reliable operation of the network. By using an OPF, the ISO can legitimately allocate generation in an "optimal" way while respecting line flow constraints, voltage magnitude constraints, VAr constraints and any other constraints that are necessary to ensure safety and reliability. As a by-product, the OPF also produces the shadow prices associated with locationally based marginal pricing (LBMP) of power. These prices can be used as determined by the market mechanism being employed.

In the context of a market in POWERWEB, the OPF may be subjected to widely varying costs and therefore dispatches which are far from typical base case operation. It is important in such an environment that the OPF be extremely robust. The latest release version of the Matlab OPF solvers used in POWERWEB and more detailed documentation of the algorithms employed are available at no cost at <http://www.pserc.cornell.edu/matpower/> as part of the MATPOWER package

3.3 Summary

PowerWeb has so far been used to examine the effectiveness of day-ahead electricity markets. Over 100 people have participated in electricity market experiments using this software platform. It has allowed for simple variation in the market mechanism being examined and also variation in the type of generators in the market. The two most important series of experiments conducted so far have examined the ability of generators to sustain prices above marginal cost in the presence of network constraints and the ability of generators to self-commit when faced with start-up costs.

4.0 Market Power Experiments

Market power increases as sellers own a larger fraction of the capacity available for serving demand. In an electric power grid, the supply and demand are dispersed throughout the system. Each generator and each load lie at a specific network location. Due to the constraints imposed by the need to operate the transmission grid reliably and securely, it may not always be possible to transfer power from an arbitrary generating station to any given load. This implies that the capacity available to serve a specific load may be a subset of the total generation capacity in the system and that market power may be present if a small number of sellers own a large fraction of this subset of generation. The market is partitioned into smaller market islands by the limitations on transmission imposed by the network. If areas A and B of a transmission grid are isolated by transmission constraint, then generator A in area A cannot compete with generator B in area B to serve the load in area B. Likewise, generator B cannot compete with generator A to serve load in area A. The owner of a generation facility may have market power if they own a significant percentage of capacity in an isolated area even if they own only a small fraction of the total generation in the system.

These transmission limits may be simple and relatively constant thermal limits on the lines or they may arise indirectly from voltage or stability limits. In the latter case, the constraints may be very sensitive to VAr (reactive power) injections and other operating conditions. Therefore, market power could also arise from ones ability to manipulate the operating condition of the network in order to partition the markets to one's own advantage. For example, consider a network with a key transmission line connecting bus 1 in area A to bus 2 in area B. And suppose that the amount of power which can be transferred from A to B (while satisfying voltage limits) is highly dependent on the VAr injection at bus 2. It may be possible for a generator at bus 2 to isolate itself from competition from area A by withholding VAr capacity.

In summary, there are at least two ways in which the transmission network can create market power opportunities in load pockets. First, transmission constraints, arising from line limits, voltage limits, or stability limits, may partition the market into islands which may create the type of market power described above. Second, one may exploit one's position in the network to strategically partition the market to one's own advantage. The simple auctions tested above do not take into account transmission system constraints. The dispatch schedule produced by such simple auctions would often lead to infeasible operating conditions if employed in a constrained network (see for example, Hogan, 1992). The answer to this problem, of course, is use of a smart market which employs an auction where offers are adjusted for nodal pricing through transmission charges determined by an optimal power flow (McCabe, Rassenti et al. 1991).

We have conducted three experiments using experienced subjects (who had participated in the LAO sessions described above) in a smart market network environment. These experiments used a LAO auction with prices and offers adjusted for location in the network via an OPF (optimal power flow).

4.1 A Smart Market

A smart market is needed to account for the operational constraints imposed by the physical transmission network. In this context, the sellers and the buyer's demands are connected by a transmission network which must be operated at all times in a manner consistent with the laws of physics governing the flow of electricity. The operation of the network is also constrained by the physical limitations of the equipment used to generate and transmit the power. This results in two phenomena that may affect the auction: (1) transmission losses and (2) congestion.

The transmission lines dissipate a small percentage of the energy produced by the generators. The amount of power lost depends on the flow in the line and the length of the line, among other things. Transmission loss implies that the total amount of power the buyer must purchase is slightly greater than the total demand and the exact amount is dependent on where the power is produced.

There are limits on the amount of electric power that can be transmitted from any given location to any other location. Some of the limits are simple line capacity limits and others are more subtle system constraints arising from voltage or stability limits. Congestion occurs when one or more of these network limits is reached. Congestion implies that some inexpensive generation may be unusable due to its location, making it necessary to utilize a more expensive unit in different location.

The effects of losses and transmission system constraints are handled by adjusting all offers and prices by a location specific transmission charge which represents the cost of transporting the electricity from the respective generating station to some arbitrary reference location. There is a two-part transmission charge associated with each line that is divided up between the various generators based on their individual contributions to the flow in the line. The per-line transmission charges can be explained as follows. The value of the power dissipated by a transmission line is the loss component of the transmission charge for that line. The congestion component of the transmission charge is precisely the charge necessary to discourage overuse of the line. If there is no congestion, this component is zero. It is important to note that the transmission charges are dependent on the flow in each transmission line as well as each generator's contribution to that flow and therefore cannot be computed before performing the auction. In this context, each generator receives a price that is specific to its location.

Units are chosen so as to satisfy demand in the least expensive manner while satisfying the operational constraints of the transmission system. An optimal power flow program that computes the appropriate transmission charges for each generating station does this. The units selected by the optimization program are roughly those given by the following procedure. The appropriate transmission charge is added to the price of each offer, and the offers are ordered from lowest to highest adjusted offer price. Units are included for sale, starting from the low priced units and moving toward the higher priced units, until the supply reaches the total buyer's demand plus transmission losses. The remaining, higher priced, units are excluded from sale.

The reigning price is set to the adjusted offer price of the last (most expensive) unit chosen. The price paid for each unit produced by a given generator is the reigning price minus the corresponding transmission charge.

4.2 Experimentation

We conducted three experiments using the POWERWEB experimental platform which implements the smart market described above using an OPF that models a full non-linear lossy AC transmission network. These experiments utilized the six generator, 30-node network model, shown as a simplified block diagram in Figure 4.

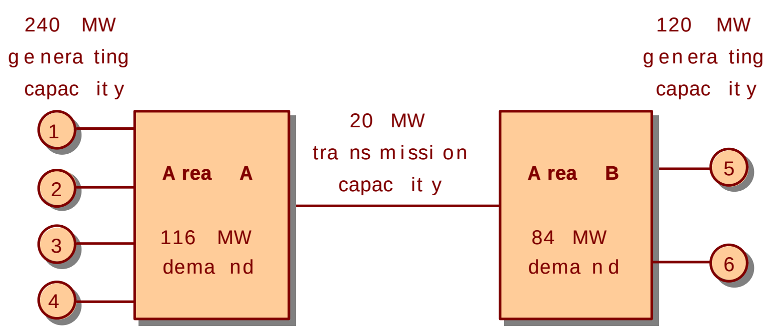

Figure 4: Transmission Network Block Diagram

Each of the six subjects in each experiment was one of six sellers in a market with a single buyer with a fixed demand. All generators had a capacity of 60 MW (megawatts) which was divided into 3 blocks, 12, 24, and 24 MW at prices of \$20, \$40, and \$50/MW-hr, respectively. All generators had identical capacity and cost structures. Each generator could generate between 12 and 60 MW of power, or could be shut down completely, in which case it incurred no costs.

The network was structured so as to create a load pocket in Area B, where generators 5 and 6 are located. The limitation on transmission capacity between areas A and B, can effectively separate

the market into groups of four and two competitors, respectively. The demand levels and network constraints are such that neither generator 5 nor generator 6 can be shut down.
Each of the three sessions was run for 75 rounds, and each produced different results. Figure 4 shows the price results for a session that can be used to characterize all three sessions. In one session, the results for the prices received by the six generators remained similar to the price pattern shown in the figure prior to period 50. In other words, prices remained near the competitive level (shown by the heavy horizontal line in the figure) throughout the session. In a second session, prices were similar to those shown after trading period 50 in the figure, for the entire session. In other words, generators 5 and 6 were able to exploit their market power consistently from the initial trading periods through period 75. In the session shown in the figure, generators 5 and 6 were not able to coordinate their price offers to exploit the market power opportunities offered by the network until period 50. It appears that generator 5 (dashed/dotted line, 2nd from top) was not responsive to generator 6 (solid line, top) who attempted to raise prices earlier.

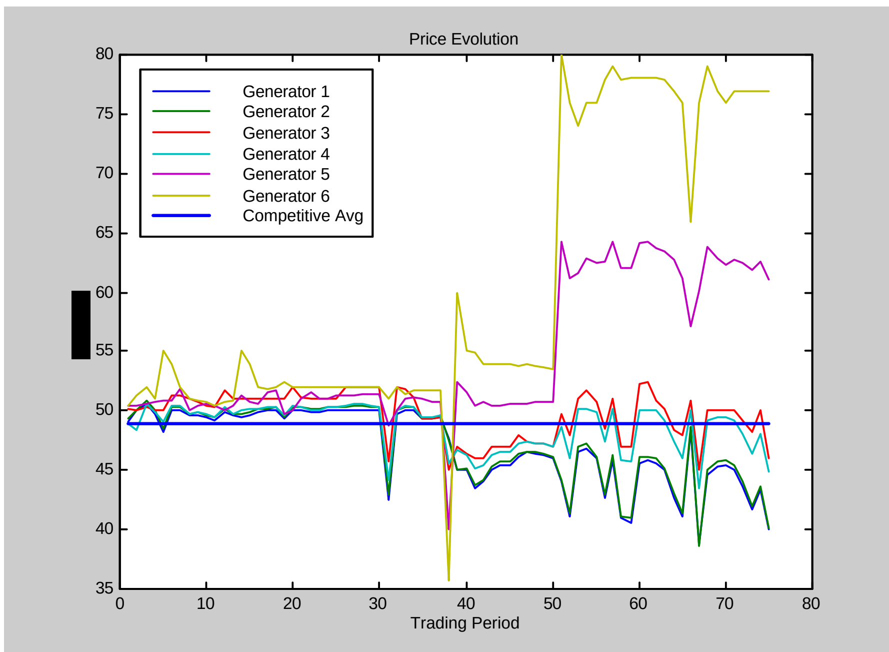

Figure 4: Nodal Prices (Students)

We draw two conclusions from these results. First, in two of the three sessions generators 5 and 6 were able to exploit the opportunity to use market power. It should be noted that the 75 trading periods used provides far less experience than actual generators will accumulate over a summer season during peak load periods when networks are likely to be constrained. Thus, it is reasonable to conclude that market power will be exercised. Second, if generators exploit market power, prices will not only be higher in load pockets, but also price volatility will increase. This implies the possibility that network stability and reliability may be jeopardized since relays have been set on the basis of stable generation patterns throughout the networks.

A nearly identical experiment with 65 trading periods was later performed using electricity traders as subjects. Once again, the market power opportunities were quickly recognized and exploited. Prices well above competitive levels were observed at generators 5 and 6 as early as the second trading period, and remained consistently high after about 25 periods. This result supports the conjecture that the behavior of expert subjects does not differ significantly from that of the more accessible student subjects. Figure 5 shows the nodal price progression from that experiment.

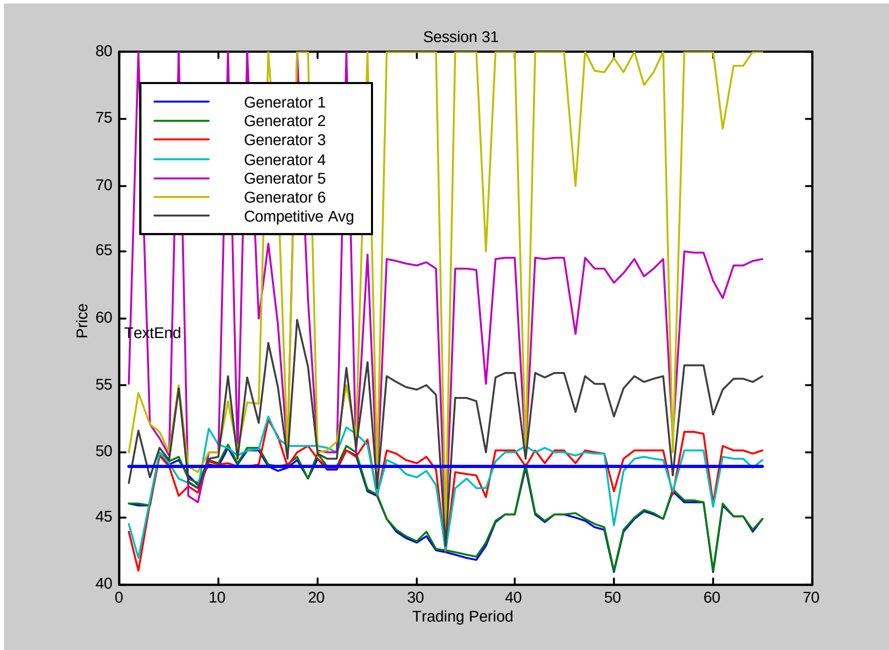

Figure 5: Nodal Prices (Industry Professionals)

5.0 Unit Commitment Experiments

5.1 Overview

Given the load profile of any electricity market and the capabilities of the generators supplying power to that market, it is likely that only a subset of the total number of generators will be required to satisfy the load. One of the most important roles for a system operator, whether it is a utility with a portfolio of generating assets, a state controlled government agency, or an independent system operator, is to determine which generators should be running and for how long. This is frequently called the "unit commitment problem". While the task has to be solved, the method can dramatically vary in different markets. For example, the system operator in the United Kingdom solves both the unit commitment and dispatch problems in a day-ahead market. In contrast, the markets in Australia and California are based on self-commitment by generators and the system operator determines dispatch only. The emerging markets in the eastern United States are closer to the United Kingdom model than to the simpler markets with self-commitment. The basic question posed by these different markets is which approach is the best?

In regulated markets the system operator often has perfect information about heat-rate curves, generator operating constraints and network constraints, and can determine the optimal dispatch schedule of thermal units. The aim has been to find the optimal production level for each generating asset that minimizes the total cost of meeting the demand for power over a specified period of time. Two main approaches have been used. The first, primal approach, seeks a series of dispatch decisions that minimize production cost. The second, dual approach seeks to find a set of prices that induces the optimal commitment of all resources. Bertsekas (Bertsekas, Lauer et al. 1983) found that not only did the dual approach lead to solutions close to those produced by the primal approach but the gap, between the two solutions, decreased as the system size increased. The complexity of the problem, however, exposes us to the possibility that in each case there may be more than one optimal or close to optimal solution. Johnson, Oren and Svododa (Johnson, Oren et al. 1996) highlight the issue that the objective total cost function may have a "flat bottom". There may be many near-optimal solutions to the unit commitment problem. They found that while the different solutions resulted in small variations in overall cost, the profitability of different generating assets varied greatly, particularly for marginal units. When ownership of units was centralized, such as under one utility company or a government agency, then this flaw is of little importance. Ownership in today's electricity industry, however, is characterized by increasing decentralization. Who should decide which of the many possible unit commitment solutions is to be implemented? For owners of more marginal, high cost generators, this could be the difference between profit and loss. For the system operator, there is no economic rationale for picking between any one of the possible solutions. This casts significant doubt on the continued feasibility of centralized dispatch in a deregulated environment.

For this reason, markets that have deregulated have turned to the use of auctions, in which power producers offer to supply power at a given price. The ultimate solution to the unit commitment problem will have been determined (at least in part) by price schedules submitted by owners of generators. Equity is restored. The auction is not without its problems. The market may be more brutally fair but can it deliver cost efficient dispatch? Oren and Elmaghraby (Elmaghraby and Oren 1999) develop the idea that by deviating from a marginal cost offer strategy some generators can significantly increase their own profitability, at the expense of system wide cost-efficiency. In particular, they posit that start-up costs can be an incentive to distort the efficiency of dispatch by providing inefficient generators an incentive to undercut the offers of more efficient generators and sneak into the dispatch schedule.

The inter-temporal dependencies caused by start-up costs provide an incentive to accept losses or reduced profits in some periods in order to increase profits overall. This may cause generators to offer blocks of capacity at below cost in order to avoid a greater start-up cost in a future period. Every generator must determine whether this increases profitability or whether its cost structure is such that it should cycle on and off with the variations in demand. In an intensely competitive market, the optimal strategy should be one that leaves the generator at worst indifferent between the cycling and continual operation. In such a case, the losses incurred would exactly equal the start-up costs avoided. For that reason, this strategy would appear consistent with aims of profit maximization. It is this basic assumption that has been tested in our research.

5.2 The Experiments

We conducted eight experiments to test this hypothesis with our web-based PowerWeb platform, which implements the smart market, described previously, using an OPF that models a full non-linear lossy AC transmission network. These experiments used a six generator, 30 node network model. Each of the six subjects in each of the experiments was one of six sellers in a market with a single buyer with demand that alternated between 100 MW and 200 MW. All generators had a capacity of 60MW that was divided into three blocks, the size of which varied between generators. The costs for each block of capacity varied between generators too. Subjects knew their own capacities and costs but not those of their competitors. Table 1. below shows the capacity and cost structure of each of the competitors:

	Variable Costs					
	Block 1		*Block 2*		*Block 3*	
Generator	**MW**	**Cost ($)**	**MW**	**Cost ($)**	**MW**	**Cost ($)**
1	10	23	25	30	25	35
2	10	23	25	30	25	35
3	20	18	30	18	10	40
4	20	20	20	30	20	40
5	20	20	20	30	20	40
6	20	15	30	15	10	40

Table 1: Generator Capacities and Variable Costs

Each generator was required to sell at least its first block of capacity in its entirety. If this did not happen, the generator was shut down for that period. In the event of being shut down, a start-up cost was incurred when the generator again was selected to operate. Table 2. shows the start-up costs for each of the generators:

Generator	*Type*	*Start-Up Cost ($)*
1	Peaking	50
2	Peaking	50
3	Base-load	500
4	Mid-Level	150
5	Mid-Level	150
6	Base-load	500

Table 2: Start-Up Costs

The network was structured so as to eliminate any network constraints. Losses in the system still occurred but were too insignificant to affect the optimal offer strategy of each generator.

Six sessions were run with undergraduate business and economics students at Cornell University. The majority of students were sophomores and juniors taking an intermediate microeconomics class and/or a class in price analysis. One experiment was run with Graduate students in economics and a final experiment was run with power industry professionals. The six undergraduate sessions and one professional session were run for 60 rounds. In 30 rounds demand was 100MW and in 30 rounds the demand was 200MW. The graduate experiment ran for 40 rounds, being also evenly split between high and low demand periods. A low demand period was always followed by a high demand period and vice versa.

A uniform price auction was held in advance of each of the trading periods. Subjects were informed of the demand for that period and asked to submit offers for each of their blocks of capacity. Units were chosen based of their offers into the auction so as to satisfy demand in the least cost manner while satisfying the constraints of the transmission system (in this experiment to include losses only). Upon submission of offers and completion of the OPF, students were presented the results and profits (based on the reported clearing price and the quantity of electricity sold in the auction) from the previous trading period before submitting offers for the next period. Subjects were paid based on their performance, experiment results were in experimental dollars. An exchange rate was applied to this and students were shown at each stage their earnings in actual dollars. Each subject received an initial "turn-up" fee, which was used as an incentive to encourage people to attend the experiment. It was then considered as a starting balance in the experiment. It was possible for subjects to lose money as well as make profits. Losses were capped at $0 (after application of the turn-up fee). There was no cap on the profits that could be made.

Our hypothesis has been that some generators would find it profitable to offer sufficient capacity so as to be dispatched at below marginal cost in order to avoid start-up costs in the next period. Invariably, given the demand and supply structure in these experiments, everyone sold something in high demand periods. The low demand periods are, therefore, of most interest. Table 3 below shows the appropriate offer strategy for each generator. The offer strategy is calculated using the following formula, applicable to two period games[1]:

On capacity < minimum capacity,
offer = average cost of block[2] *- start-up cost/ size of first block*
On capacity > minimum capacity,
offer = marginal cost

	Block 1		***Block 2***		***Block 3***	
	Cost ($)	**Offer ($)**	**Cost ($)**	**Offer ($)**	**Cost ($)**	**Offer ($)**
1	23	18	30	30	35	35
2	23	18	30	30	35	35
3	18	-7	18	18	40	40
4	20	12.5	30	30	40	40
5	20	12.5	30	30	40	40
6	15	-10	15	15	40	40

Table 3: Optimal Offers

5.3 The Results

The experiments validated the hypothesis that last accepted offer auctions can produce cost efficient dispatch. The graphs on the following page show the offer strategy of each of the six

[1] If all generators followed this strategy, optimal dispatch of generators would occur.
[2] Because each MW in a block is the same price, average cost equals marginal cost. It is appropriate, however, to think in average cost terms because in the US power auctions often restrict the number of segments in a price/offer schedule. This forces generators to offer blocks of capacity at the same price.

generators averaged over all of the undergraduate sessions in low demand periods. The upper boundary straight line is the offer expected if the generator submitted only marginal cost offers. The lower boundary represents the offer predicted which would leave the generator indifferent between being on in both periods or being on only in high demand periods. In reality the cost structures of the generators in the experiments meant that different generators faced different degrees of competition. The baseload generators faced the least competition while competition was fiercest between generators 1,2 (ordinarily cycling) and generators 4,5 (ordinarily dispatched). We believe that an offer pattern between marginal cost and lowest possible offer can be considered (close to) optimal.

As expected, generators 1,2 and 4,5 all converge on the predicted offer. The base-load generators were under less competitive pressure. Nonetheless, their offers also sank below cost in low demand periods, though to a lesser extent. This merely reflects the fact that it is only rational to lower the offer until dispatch is secured. For the base-load generators in this experiment, that was significantly higher than the minimum offer suggested in this paper. These results were also replicated in the graduate and professional experiments.

Figure 6. shows the cost efficiencies of the experiments over cycles of one high and low periods. It's a messy picture but one which conveys the convergence of each of the experiments to close to 100%. Efficiency in these experiments is defined as optimal system cost divided by realized system cost. By means of comparison, had generators submitted marginal cost offers, the efficiency would have been just over 96%. The results show that self-commitment using a uniform price auction converged to a higher efficiency than this.

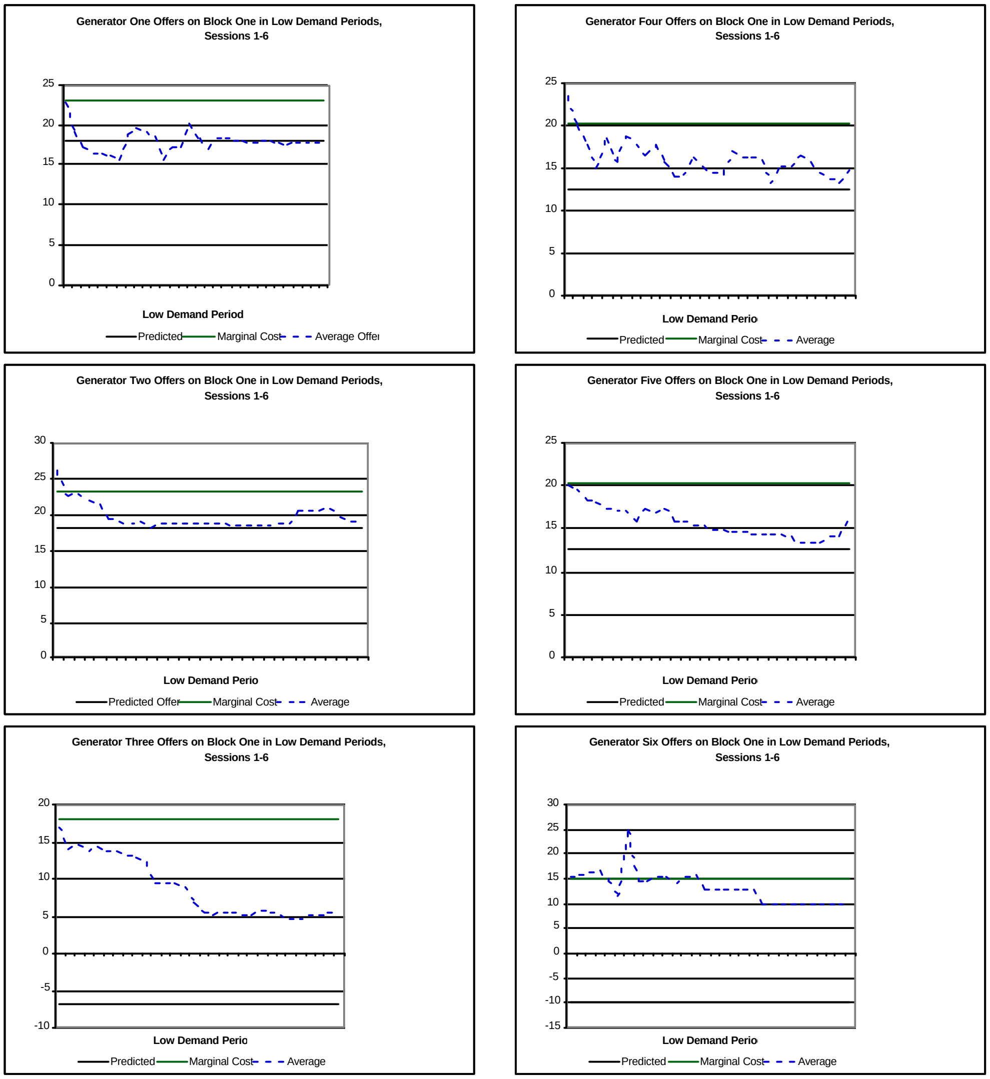

Figure 5: Low Demand Period Offers in Undergraduate Sessions

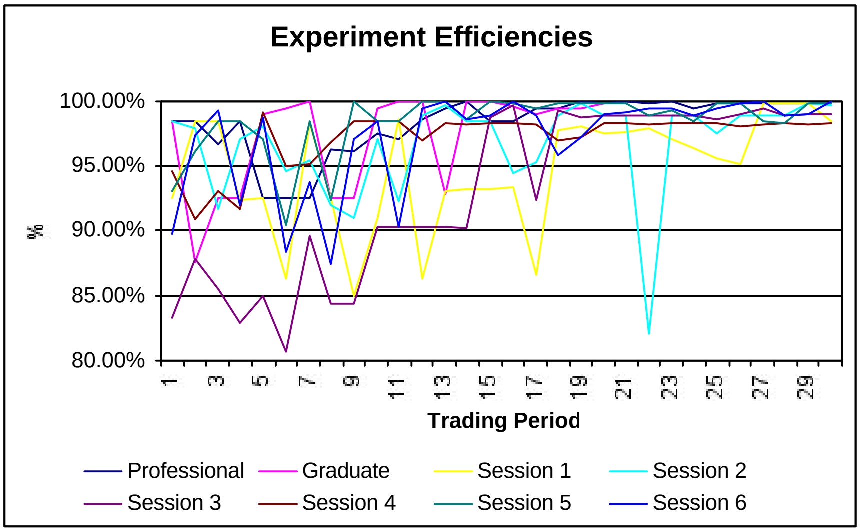

Figure 6: Experiment Efficiencies

Figure 7 shows the average efficiency of the undergraduate experiments compared to the efficiency of the graduate and the professional experiments. The only difference that can be seen between the three groups is the speed with which optimal dispatch was achieved. This again supports the conjecture that behavior of expert subjects does not differ greatly from more accessible student subjects.

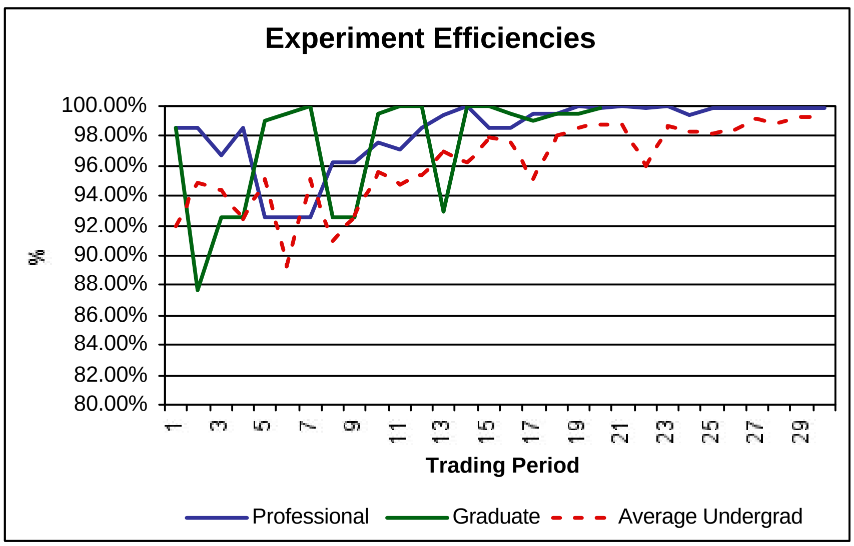

Figure 7: Comparison of Efficiencies

Our experiments show, in a simplified situation, self-commitment can produce a cost efficient dispatch of thermal units. Further complexity needs to be added to the model in the form of ramping constraint and minimum up and down times before it is possible to conclusively say that self-commitment is feasible. Nonetheless, the success of the uniform price auction in this instance is encouraging, given its position of auction-of-choice in electricity markets. Had it failed this simple test, severe doubt would be cast upon its ability to handle more complicated scenarios.

6.0 Conclusion

The experiments using POWERWEB have shown that day-ahead markets, in the absence of transmission constraints produce cost efficient outcomes. Low demand period prices can be much lower than cost as a result of some generators attempts to avoid de-commitment and start-up costs. Transmission line constraints, however, can produce sustainable high prices in isolated regions in the network and by cascading effect throughout the network. Which of these two effects dominates will depend upon the network structure and the cost structure of the generators supplying power to the system. With this basic understanding of the energy market (or the market based forecasted load) it is now possible begin an examination of the ancillary services market.

7.0 Alternative Markets for Reserves

Reserve markets are a method of hedging the risk that pre-determined supply of electrical power from a forward energy market will be insufficient to meet actual demand in real time. Having used PowerWeb to simulate an electricity market with randomized demand and load forecasting errors, we examine different means to mitigate both price spikes that exist in electricity markets and the concern to ensure adequate real-time supply of electricity at least cost.

We are examine three possible approaches to the problem of securing reserve generation and each method's effect upon price volatility. Those approaches can be summarized as follows. Firstly, both energy and reserve markets can be settled real-time. Secondly, the energy market can be settled ex-ante based on forecast load with a balancing market clearing the different between actual and forecast load in real time. The former method is similar to that used by New England Pool. The latter approximates the proposed new PJM market structure. A third way is to settle both energy and reserve markets ex-ante and run a real time balancing market for both energy and reserves. The method is not used in any real markets but possesses some favorable characteristics worth investigating.

This section will describe in greater detail each of these three market structures and hypothesize their comparative effects upon price volatility. Results from a pilot experiment, using POWERWEB, examining price spikes and the proposed PJM market structure will also be presented.

7.1 Real Time Settlement (New England)

In New England (NEPOOL) energy and reserve markets are settled in real-time when actual load is known. Load is forecast a day ahead. At this point, generators submit binding offers for their available capacity into the forward energy and reserve markets. The independent system operator (ISO) simultaneously optimizes for energy and reserve markets. In the optimizing algorithm, greater weight is placed on minimizing cost in the energy market than in the reserve market. Generators are subsequently given a non-binding indication of whether they will be dispatched in the energy market or be held as reserves. The expected prices are calculated for the energy and reserve markets markets. Although the markets are not settled until real-time, when cleared and settled they are done so on the basis of the offers submitted a day ahead. Those generators dispatched in the real-time energy clearing market are paid the real-time energy-clearing price. Those generators held as the reserves receive the lower of the energy clearing price or the reserve market price. This is in effect and ex-post payment for having provided available capacity. Figure 1, shows a schematic of this process.

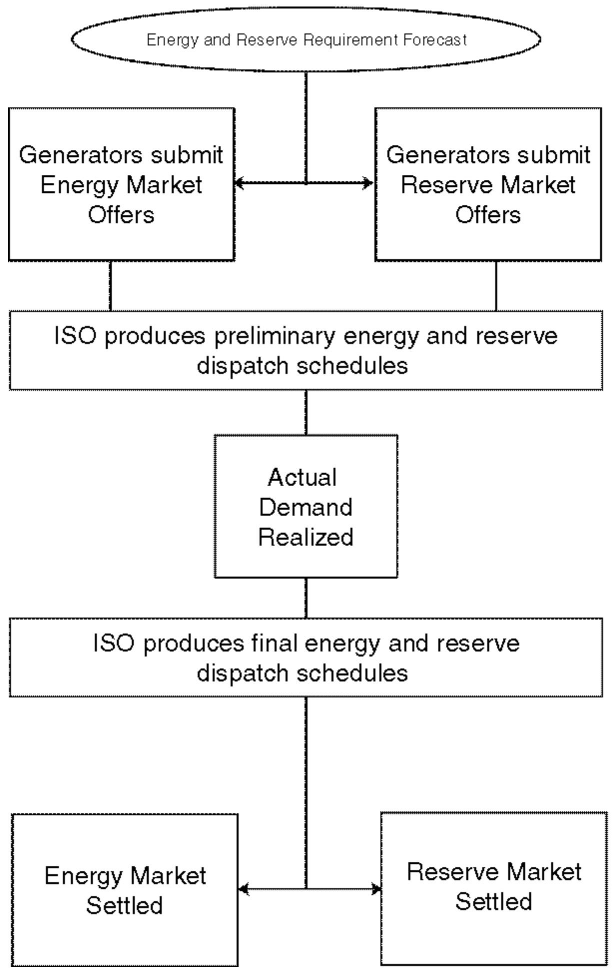

Figure 1: Schematic of Real-Time Energy and Reserve Market

There are two important ways to compare the operation of different reserve markets. Firstly, do they provide sufficient incentives for power producers to provide reserve capacity? This is particularly important when producers are able to supply power to more than one network. Failure to be dispatched sufficiently in one market may lead to switching of the remaining capacity to an alternate market. If this is prevalent, sufficient reserve capacity may not exist. Perhaps more importantly than actual level of remuneration, in this respect, is the certainty of remuneration.

A second criteria for a reserve market is its ability to mitigate the effect of price spikes upon the cost of supplying power. The bulk of load can be predicted with almost certainty. It is the variation in the tails of the load's distribution that causes forecast to err from actual load. This variation in the tail of the distribution also produces the price variability. If the reserve market is intended to compensate for variability in load, does it reduce or increase the cost of meeting the extra load.

The real-time energy and reserve markets do provide an incentive to supply reserve capacity. The indicative statement from the ISO informs generators of their expected schedule, the expected price and expected return from participation in the market. The binding offers in the forward markets ensures that there will be sufficient reserve capacity in the real-time markets. The main drawback to this system is that the effect of price spikes is entirely focussed upon the real-time market. If forecast demand underestimates actual demand the price of energy increases for the entire load, not just the extra portion required.

7.2 Balancing Market (PJM Proposal)

A second approach employs a "balancing market" in real-time to settle differences between the forward energy market outcome and the real-time energy requirements. Generators submit offers into a forward energy market based upon forecast demand for load. The forward market is solved and settled and a forward energy-clearing price is determined. Generators are then permitted to submit offers into a balancing market for that capacity not dispatched in the forward energy market. If actual demand is greater than forecast demand, the balancing market is cleared to meet the extra demand. A balancing market price is paid on the extra capacity dispatched in the balancing market. If actual demand is less than forecast demand, the forward energy market is resolved. A new lower energy-clearing price is determined with a smaller dispatch schedule. Those generators backed down pay back on the capacity dropped at the new energy-clearing price. The margin between the old and new price is a per-unit payment for having provided available capacity. This system is similar to the proposed market format for the PJM system.

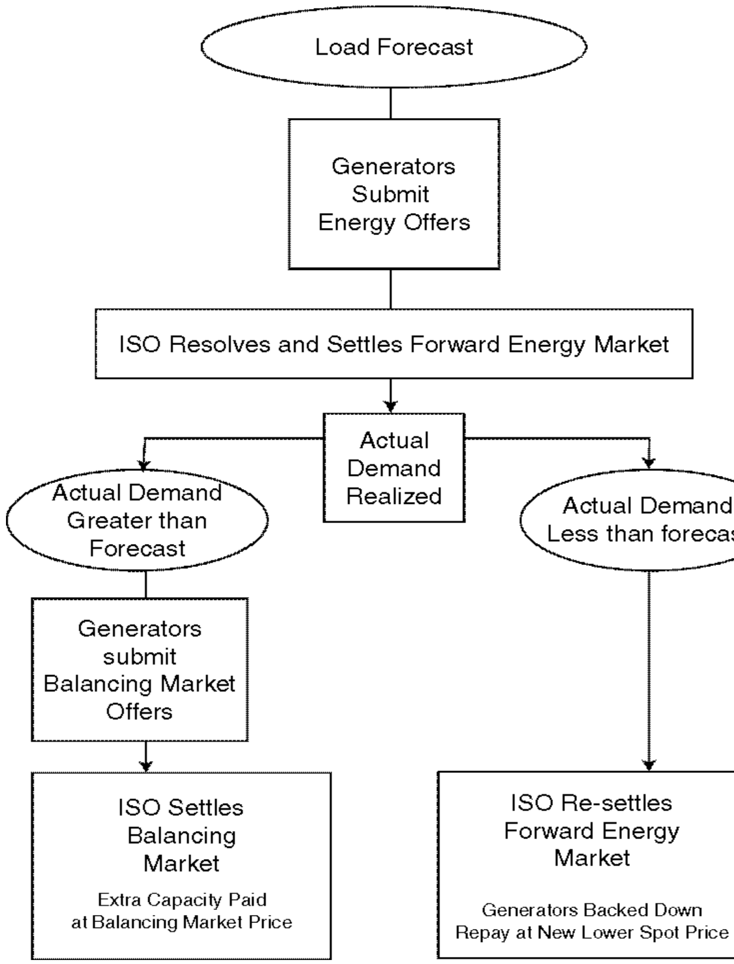

Figure 2: Schematic of Energy Market with Balancing Market

Figure 2 shows a schematic of the balancing market system. There is no reserve market in this system. A reserve margin can be built into the forecast load but there is no formal market mechanism. The reserve margin would have to be settled at the energy clearing price. Alternatively, the reserve capacity could be mandated by the system operator and reimbursed using a regulatory formula.

Since profits from expected load are locked in, the balancing market is effectively a second chance to earn revenue on capacity that was not dispatched in the forward market. The incentive for generators to participate in the balancing market is the anticipation that profits will be high when it is employed. Even if dispatch is not certain, expected profits may be sufficient to induce participation (and by default act as reserve capacity.)

By settling the forecasted load and forward market ex-ante, the extra-portion of load, alone, must bear the higher balancing market price. The effect of a price spike is, therefore, isolated from the bulk of the load. Those customers who demand more power than anticipated must incur the entire cost of meeting the extra supply. This compares to the real-time market where the cost of meeting the extra load is mutualized across all consumers.

7.3 TIM Market

The final method operated a forward energy and reserve market, into which generators simultaneously submit offers for their available capacity. Both markets are cleared and settled, deriving forward energy-clearing and reserve prices. If actual demand is greater than forecast demand then the forward market is resolved and extra capacity is dispatched in order of offers (dependent upon being in reserve market). The extra capacity is paid the aggregate of the forward energy price and the reserve price. If actual demand is less than forecast, generating capacity is backed off and paid the forward reserve price for the capacity dropped and refunds to the ISO the forward market revenue for that capacity.

Figure 3 shows a schematic of the TIM market. Like the balancing market, the cost of meeting additional demand is borne by those causing that extra load. This reduces the average cost of meeting actual demand. This extra cost will be less than with the PJM market since the energy clearing price is not resolved. In addition, the market has the advantage that it also secures reserve capacity in advance. For some generators this may not be sufficient to recover costs. The reserve market offer should account for this, so that if the generator is called upon to generate it does not lose money.

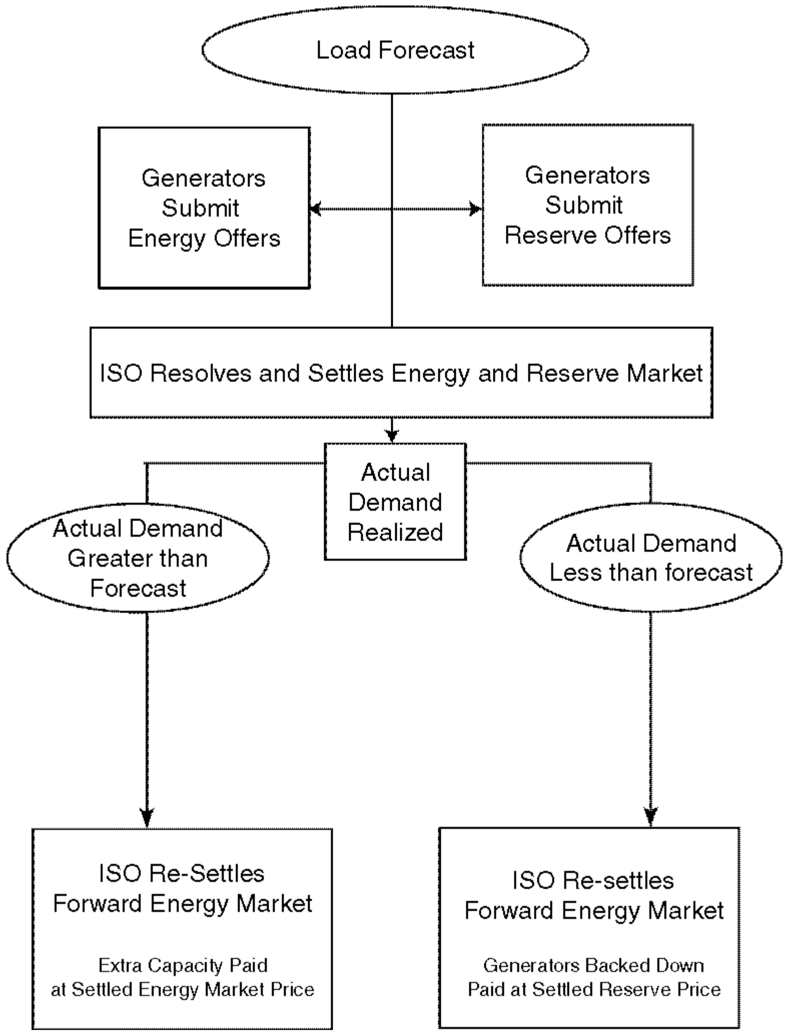

Figure 3: Schematic of TIM Market

7.4 Summary

The greater the proportion of demand settled in the forward market, isolating the additional capacity that causes the spike, the more the spikes in average prices are mitigated. The second and third models possess this characteristic. The cost of meeting actual demand for load in these two markets is substantially lower, during periods when price spikes would occur, than the NEPOOL market. This is not an unexpected result. While it is important to consider the effect upon average prices, the ability of markets to induce supply of reserves is also important. The PJM style market does not guarantee in advance the availability of reserve generation. It relies on the expected price in the reserve market providing sufficient incentive to generators to be available. The balancing market will periodically experience very large price spikes. Generators will only provide capacity to this market if expected returns exceed those from the next best alternative market for the capacity. Where it is difficult to supply power to outside of the pool then this presents less of a problem. There no guarantee that reserve prices will be lower than energy prices when generators are submit offers into forward energy and reserve markets simultaneously. The opportunity cost of being dispatched in the energy market to serve as reserve generation increases offers above the production cost of reserves. Finally, previous simulations have shown that when the number of competitors is small in an electricity market, energy offers are high compared to production costs. The number of generators participating in the reserve market will be smaller than the energy market. This increases the ability of participants to affect price.

8.0 The Pilot Experiment

When employing laboratory methods to analyze an economic problem, it is customary to first run a pilot experiment. It enables the researchers to assess the experiment mechanism for its clarity of purpose and ease of use. Most subjects who participate in economic experiments do so for the first time. It is important that they comprehend the context of the experiment and what is required to successfully participate. A pilot experiment also allows the researcher to gauge whether a full series of experiments is likely to yield conclusive results and whether it will be possible to accept or reject the hypothesis investigated.

We conducted a single pilot experiment with our POWERWEB platform to assess whether it would be possible to generate price spike inducing behavior from human subjects (as proposed earlier in this report) and to adapt POWERWEB to incorporate a reserve market. Encouraging results were achieved, though their statistical importance should not be overestimated given the nature of the experiment. We decided to test the proposed PJM balancing market mechanism in the pilot experiment.

8.1 The Experiment

The experiment used a six generator, 30 node full non-linear lossy AC network. The experiment employed three human subjects and three computer agents. Each of three subjects and three agents represented a seller in a market with a single buyer with a demand that varied between 200MW and 330MW. All of the generators had a capacity of 60MW, making total available supply 360MW. Each generator's capacity was divided into three blocks of 12,30,18 MW respectively. Table 1 below shows the capacity and costs of each of the generators. The table indicates whether a human subject or a computer agent controlled it.

Generator	**Controller**	**Block 1 Cost, $ (12MW)**	**Block 2 Cost, $ (30MW)**	**Block 3 Cost, $ (18MW)**
1	Human	30	40	45
2	Computer	30	40	45
3	Human	25	30	50
4	Computer	25	30	50
5	Human	20	33	60
6	Computer	20	33	60

Table 1: Ownership and Cost Structure of Generators

The experiment was partitioned into four different periods. In the first twenty periods, a one-stage market was operated in which forecast demand was exactly equal to actual demand. There was, consequently no balancing market operated. In the second twenty periods, actual demand was different from forecast demand. In these periods, a balancing market was run when demand exceeded the forecast. When demand was less than forecast, the forward energy market was resolved, as described in the previous section. Figure 4, below, shows the demand profile from the experiment. You will notice that after period twenty, forecast demand diverges from actual demand. Forecast demand was established so that it would be less variable than actual demand.

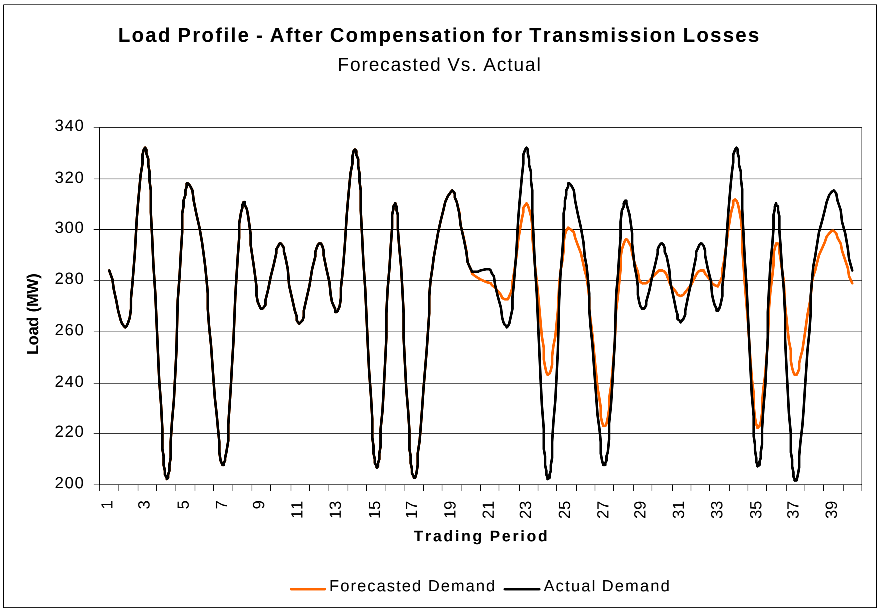

Figure 4: Forecasted and Actual Demand

Each half of the experiment was further divided into two parts. The behavior of the computer agents was different between each half. In order to introduce the human subjects to an electricity market with price spike behavior, in the first ten periods, the computer agents acted more "aggressively" than in the second ten periods. This was repeated for the second half of the experiment period. In the first ten periods, it was easier to achieve a price spike because 2 of the 3 computer agents were submitting their third block at $115/MW. In the second ten periods, only one of the computer agents submitted its third block at this higher than marginal cost. Otherwise the computer agents submitted marginal cost offers. Table 2, below, shows which of the computer agents was acting aggressively at each stage of the experiment.

Computer Agent	Trading Periods 0-10	Trading Periods 11-20	Trading Periods 21-30	Trading Periods 31-40
2	Passive	Passive	Passive	Passive
4	Aggressive	Passive	Aggressive	Passive

5	Aggressive	Aggressive	Aggressive	Aggressive

Table 2 : Behavior of Computer Agents

The human subjects were informed at the start of the experiment of the changing behavior of the computer agents. They were told that in periods 0-10 and 21-30 that the agents would act more aggressively than in the other periods. The subjects were not told the precise offer strategy of an aggressive or passive agent.

The load profile was constructed so as to involve three main peaks in demand. These were considered realistic opportunities for a price spike to occur. The peaks were such that the first would always occur when two of the agents were acting aggressively. The other peaks in demand required aggressive human subject behavior to induce a price spike. Table 3 shows for each of the four periods when the price spikes occurred.

Trading Periods	Peak Periods
0-10	3,5,8
11-20	14,16,19
21-30	23,25,28
31-40	34,36,39

Table 3: Peak Demand Periods

In the second half of the experiment a two stage market was operated. The first stage replicated the market operating in the first twenty periods, except subjects were informed that load was now only a forecast. If actual demand was greater than forecast, subjects were able to submit offers into a balancing market on capacity not sold in the forward energy market. If the generators were completely dispatched in the energy market they were sent directly to their results stage by POWERWEB. If actual demand was less than forecast, all subjects were sent directly to their results page where they were able to see how much capacity was backed off and at what cost.

8.2 Results

The most encouraging result from the experiment was replication of price spikes with human subject involvement. This suggests that in a longer experiment, human subjects may settle into a pattern of offering the third block of capacity at or near the reservation price. The least encouraging result was the similarity of prices in the forward and balancing markets. Higher prices would have been expected in the balancing market but this did not occur. The likely reason is that the experiment was too short for generators to sufficiently acclimatize to the balancing market structure. Whereas, there were forty periods in which to understand price spikes, the balancing market operated for only twenty in which only one half did demand exceed forecast to require balancing offers. Figure 5 shows the comparison of stage 1 prices against price that would have resulted from marginal cost offers from all generators.

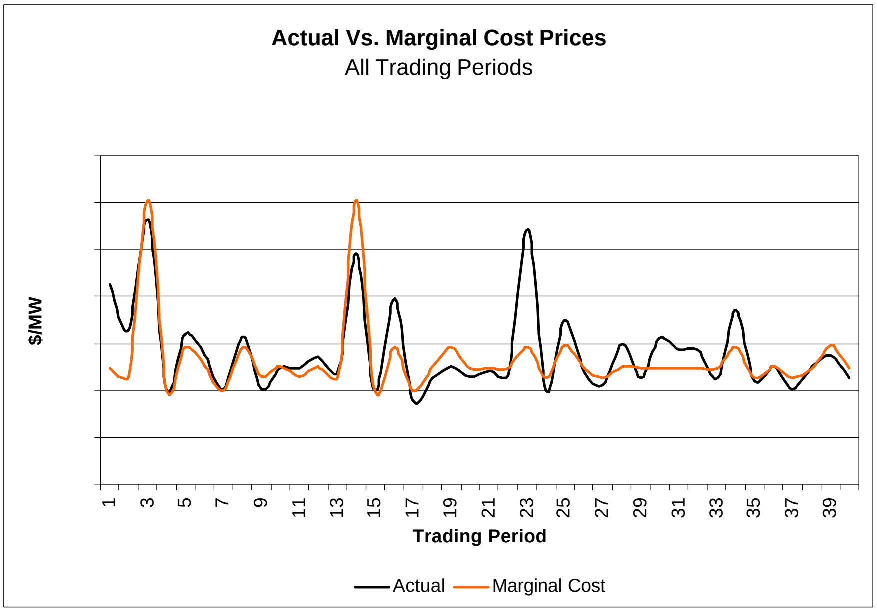

Figure 5: Comparison of Actual Stage 1 Prices and Marginal Cost Prices

It is possible to see at this level, that prices were more spikey than would have occurred with marginal cost offers. Figure 6 shows two graphs for the first and second ten periods of the experiment.

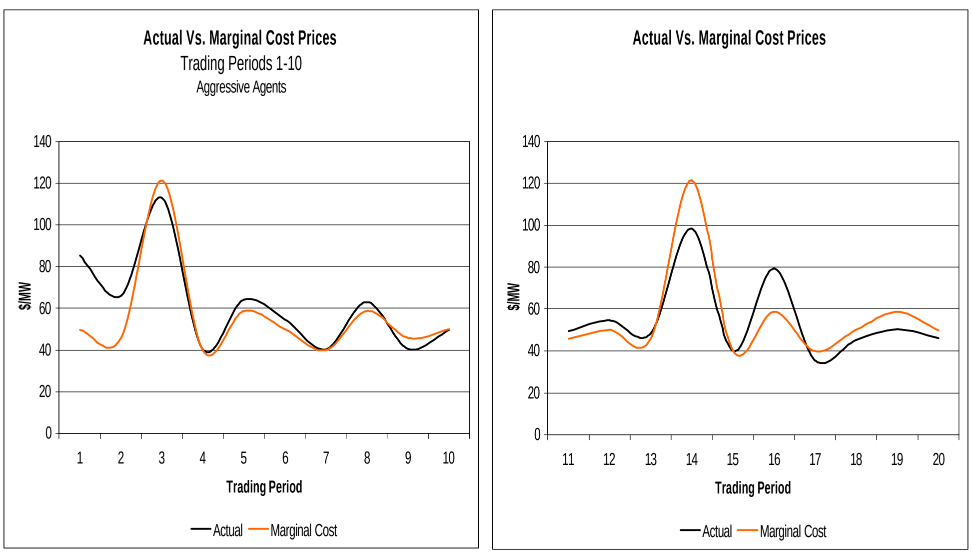

Figure 6: Comparison of Actual and Marginal Cost Stage 1 Prices, periods 1-20

The price spikes in periods 3 and 14 were expected since demand was high enough to require no human subject involvement. The first ten periods do not appear to show any aggressive offer behavior from the human subjects as the actual price follows marginal cost closely after the third tarding period. The spike in period 10, however, required an aggressive bid from one of the human subjects. Figure 7 shows the stage 1 price comparison from the second twenty trading periods.

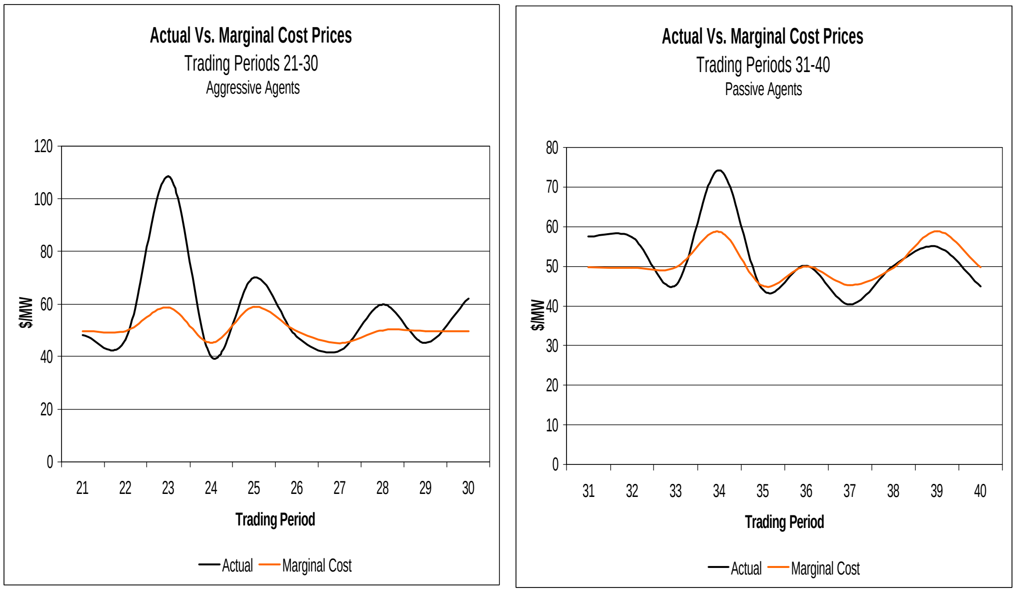

Figure 7: Comparison of Actual and Marginal Cost Stage 1 Prices, periods 1-20

The spiky price behavior in periods 25, 28 again indicate human subject involvement, though the introduction of the balancing market appears to have reduced the spiky behavior somewhat. This is likely a result of the smaller variations in load resulting from stage 1 load being an estimate as opposed to actual (as in the first twenty periods).

An examination of the third block offer strategy of the generators confirms the intuition from figures 6 and 7 that at least one of the human subjects was following an aggressive offer strategy on their third block. Figure 8 shows that generator 1 was indeed following an aggressive strategy. With a marginal cost of $45/MW the subject increased his offer substantially above cost when demand was high. The higher offers coincide with higher periods of demand.

Figure 9 shows the offer behavior of generator 3. Clearly, the offer strategy is less aggressive, though in the last twenty periods the offers were consistently above marginal cost. In figure 10, an alternative strategy was followed. Recognizing the peaks in demand, the generator submitted below cost offers to ensure complete dispatch. This is only rational if the subject was sure it was not going to set the energy clearing price. If it had done it would have lost money (this did occur but not sufficiently). Effectively, generator 5 was a free-rider on the offer strategies of the other generators.

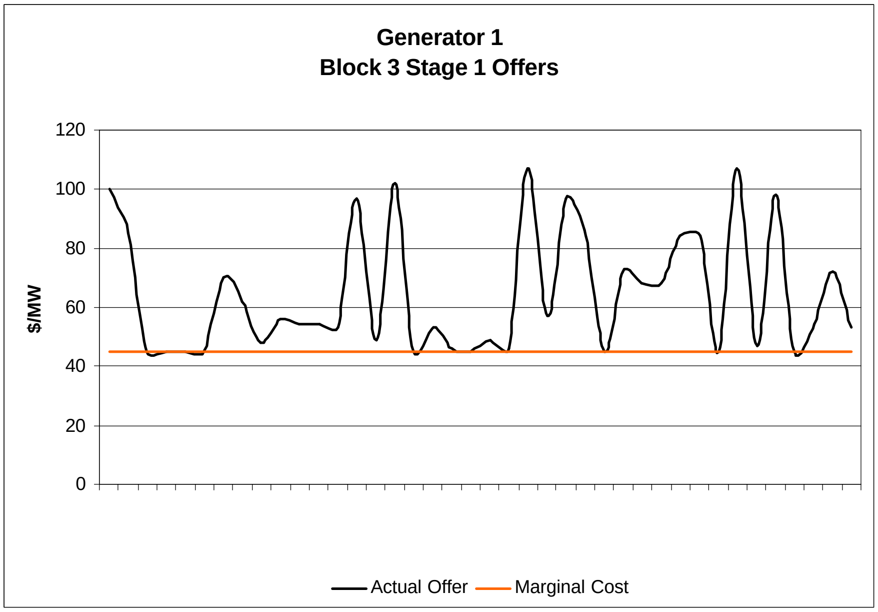

Figure 8: Generator 1, Block 3 Offers in Stage 1 Market

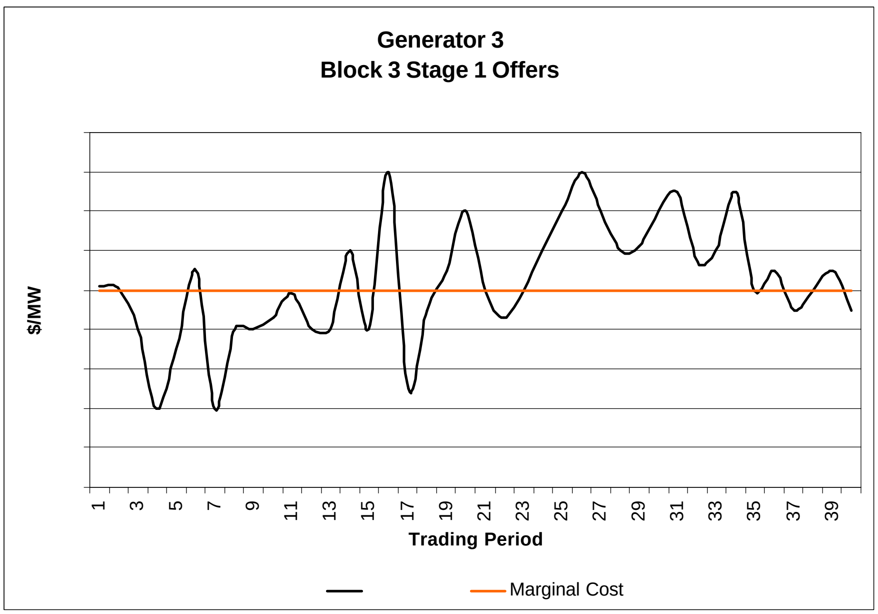

Figure 9: Generator 3, Block 3 Offers in Stage 1 Market

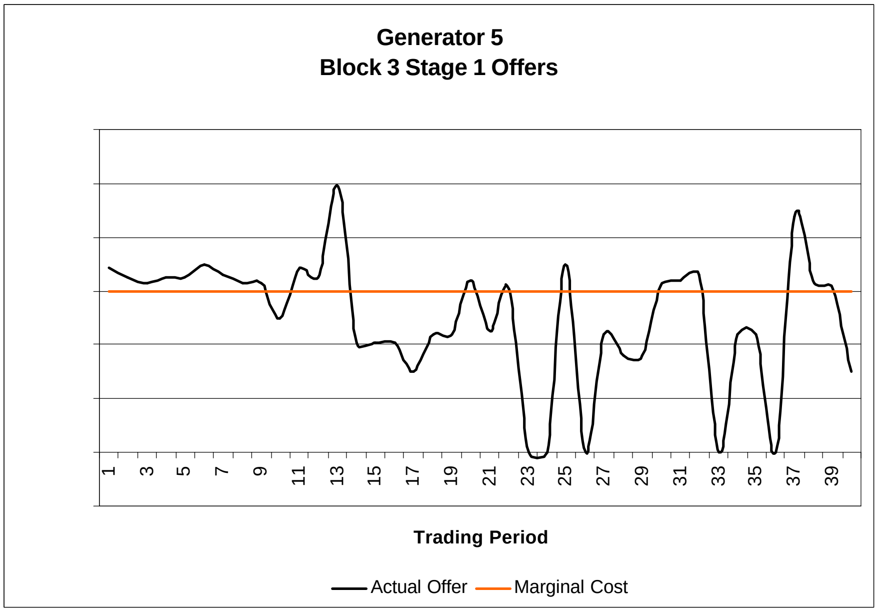

Figure 10: Generator 5, Block 3 Offers in Stage 1 Market

Figure 11 shows a comparison of the stage 1 price and the price in the balancing market. It shows that there was little difference in the prices.

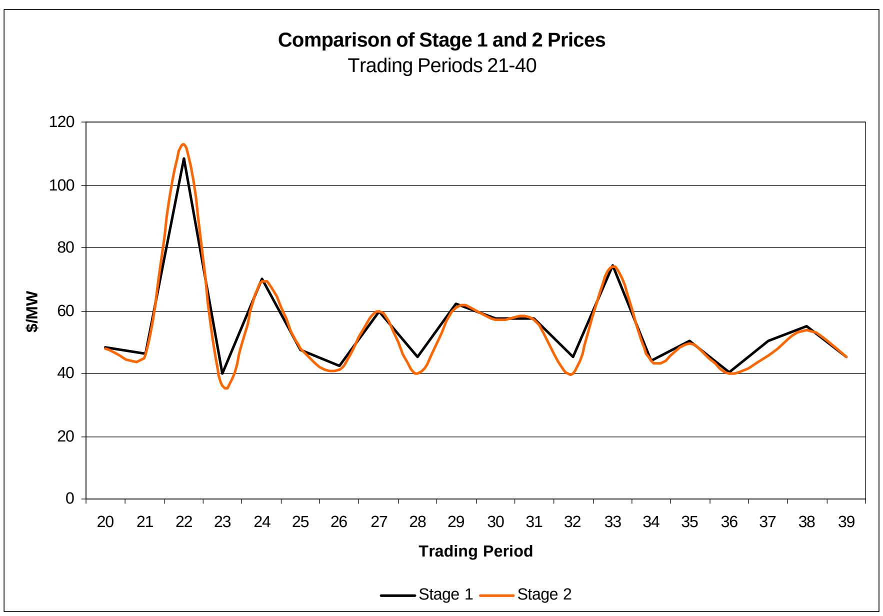

Figure 11: Comparison of Stage 1 and Balancing (Stage 2) Prices

8.4 Conclusions

POWERWEB has now been adapted to run a two stage market that will be necessary to evaluate ancillary service markets. The software's performance was robust. The pilot experiment indicates that it is possible to generate price spikes in laboratory conditions.

This will enable further analysis and experimentation. Without price spikes, which are inherent in deregulated electricity markets, it would not be possible to reach any conclusions on the operation of different reserve market mechanisms. In order to more comprehensively test the PJM market mechanism, longer experiments (perhaps over several repeat sessions) will be required. In addition it will be necessary to adapt POWERWEB further to run experiments on the NEPOOL market and the TIM market mechanisms.

9.0 Summary and Conclusions

A wholesale market for bulk power in Pennsylvania, New Jersey and Maryland (PJM) has been operating since April 1997. Prior to April 1999, the rules for the auction required that suppliers submit cost-based offers to sell power. This year the rules were modified to allow market-based offers (some suppliers have chosen to continue submitting cost-based offers). Average weekday prices for certain peak periods were higher in 1999 when market-based prices were allowed. This phenomenon could be caused by unusually hot weather or it could be caused by market power or it could be caused for other reasons. Although data on actual offers are not publicly available, it seems likely that the peak prices of well over $200/MWh were set by large suppliers with control over substantial amounts of generating capacity and not by small suppliers who control a single expensive peaking unit. The basic logic is that there is little cost to speculating with a few marginal units in a large portfolio of generators. The possibility of setting a high price on rare occasions is adequate compensation for having lower capacity factors on the marginal units. However, there are other price spikes that occur on a regular basis and could be caused by not market power but rather by the strategic posture players take relative to the market. This is a hypothesis we wanted to test.

There are several so-called ancillary services needed to ensure that an electric power network will reliably and efficiently transfer energy from generator to load. One of these services is reserve. Generators can supply both energy and spinning reserve but only in terms of a well defined relation between them. In a market (auction) setting, it is not clear what market design produces the most efficient tradeoff for a generator between offers to sell energy and offers to sell reserves and yet will provide the proper incentives for generators to want to offer reserves as a commodity. This is a key question we set out to address.

The first step in the analysis was to simulate load uncertainty and to evaluate the corresponding effects on the volatility of the market-clearing price. The objective was to identify the type of market behavior, in terms of the structures and offers submitted to the spot market for energy that leads to price spikes. This type of price volatility is characteristic of many of the new restructured markets for electricity, and it is indicative of economic and engineering problems with the performance of the spot market. An important issue for judging the performance of the different auction structures is the effect of the market structure on price volatility as well as on maintaining system security.

We were able to generate the type of price volatility observed in some real-world markets using the PowerWeb platform with computer agents representing the market participants. We used a stochastic, varying demand forecast with a relatively small random forecast error. The computer agents were programmed to submit competitive offers for those units, which have a high probability of being accepted in the market. This mimics behavior of people previously tested in real experiments. Those units, which have a very small probability of ever being chosen, if offered at cost, are instead offered at a very high price. This reflects the fact that there is very little to be lost and much to be gained by raising the price on this unit, since it is likely to be excluded anyway in all but the most exceptional circumstances. This scenario resulted in price spike behavior similar to that observed in several existing markets such as PJM. We then used this strategy as a basic bidding strategy for three integrated energy/reserve markets to be tested.

Any market for reserve generation capacity can be considered a method of hedging the risk that a supply of electrical power to meet a forecasted load from a day ahead energy market will be insufficient to meet actual demand in real time. Having used PowerWeb to simulate an electricity market with randomized demand and load forecasting errors, we examined different means to mitigate price spikes that exist in electricity markets and to mitigate the concern to ensure adequate real-time supply of electricity.

We examined three possible "reserve" market approaches. The first method was to operate a real-time market for energy. Generators submit binding offers, simultaneously, for their available capacity into the energy market and the reserve market based upon forecast load. The independent system operator (ISO) simultaneously optimizes for energy and reserve markets. Generators are subsequently given a non-binding indication of whether they will be dispatched in the energy market or are held as reserves and the relative prices for those markets. The markets are not cleared and settled, however, until actual demand for load is realized and the settlement done on the basis of the previously submitted offers. Those generators dispatched in the energy clearing market are paid the real-time energy-clearing price. Those generators held as reserve receive the lower of the energy clearing price or the reserve market price. This is in effect an ex-post payment for having provided available capacity. This formula is an approximation of the New England system.

The second approach employed a "balancing market" in real-time to settle differences between the forward energy market outcome and the real-time energy requirements. Generators submit offers into a forward energy market based upon forecast demand for load. The forward market is cleared and settled and a forward energy-clearing price is determined. Generators are then permitted to submit offers into a balancing market for that capacity not dispatched. If actual demand is greater than forecast demand, the balancing market is cleared to meet the extra demand. A balancing market price is paid on the extra capacity dispatched in the balancing market. If actual demand is less than forecast demand, the forward energy market is resolved. A new lower energy-clearing price is determined with a smaller dispatch schedule. Those generators backed down pay back on the capacity dropped at the new energy-clearing price. The margin between the old and new price is a payment for having provided available capacity. This system is similar to the proposed market format for the PJM system.

The final method of our design operated a forward energy and reserve market, into which generators simultaneously submit offers for their available capacity. Both markets are cleared and settled, deriving forward energy clearing and reserve prices. If actual demand is greater than forecast demand then the forward market is resolved and extra capacity is dispatched in order of offers (dependent upon being in reserve market). The extra capacity is paid the aggregate of the forward energy price and the reserve price. If actual demand is less than forecast, generating capacity is backed off and paid the forward reserve price for the capacity dropped and refunds to the ISO the forward market revenue for that capacity.

Several conclusions come from these simulations. The greater the proportion of demand settled in the forward market, isolating the additional capacity that causes the spike, the more the spikes in average prices are mitigated. The second and third models possess this characteristic. The cost of meeting actual demand for load in these two markets is substantially lower, during periods when price spikes would occur, than the second New England style market. This is not an unexpected result. While it is important to consider the effect upon average prices, the ability of markets to induce supply of reserves is also different. The PJM style market does not guarantee in advance the availability of reserve generation. It relies on the expected price in the reserve market providing sufficient incentive to generators to be available. The balancing market will periodically experience very large price spikes. Generators will only provide their capacity to this market if expected returns exceed those from the next best alternative market for the capacity. This presents less of a problem where it is difficult to supply power outside of the pool. There is no guarantee that reserve prices will be lower than energy prices when generators are allowed to submit offers into forward energy and reserve markets simultaneously. The opportunity cost of being displaced in the energy market to serve as reserve generation causes offers to be increased above the production cost of reserves. Finally, previous simulations have shown that when the number of competitors is small in an electricity market, energy offers are high compared to production costs. In this case the number of generators participating in the reserve market will be smaller than in the energy market which will further increase the ability of participants to influence prices.

Finally, the role of demand-side activities in competitive electricity markets was begun. The approach is based on a generalization of the competitive power pool framework to include demand-side bidding. We cast the competitive power pool into the unit commitment problem framework in which the supply cost functions are replaced by the bids submitted by the suppliers. To enhance competition, customers are allowed to play a proactive role in the price determination process by submitting bids of load reduction in specific periods.